The DOE Science Grid

Computing and Data Infrastructure for Large-Scale Science

William Johnston,
Lawrence Berkeley National Lab
Ray Bair, Pacific Northwest National Lab
Ian Foster, Argonne National Lab
Al Giest, Oak Ridge National Lab
Bill Kramer, National Energy Research Scientific Computing Center
and the DOE Science Grid Engineering Team

http://doesciencegrid.org

➢The Need for Science Grids

- The nature of how large scale science is done is changing
  - *distributed data, computing, people, instruments*
  - *instruments integrated with large-scale computing*
- *"Grids" are middleware designed to facilitate the routine interactions of all of these resources* in order to support widely distributed, multi-institutional science and engineering.

➢Distributed Science Example: Supernova Cosmology

- "Supernova cosmology" is cosmology that is based on finding and observing special types of supernova during the few weeks of their observable life
- It has lead to some remarkable science (Science magazine's "Breakthrough of the year award" (1998): *Supernova cosmology indicates universe expands forever*), however it is rapidly becoming limited by the ability of the researchers to manage the complex data-computing-instrument interactions

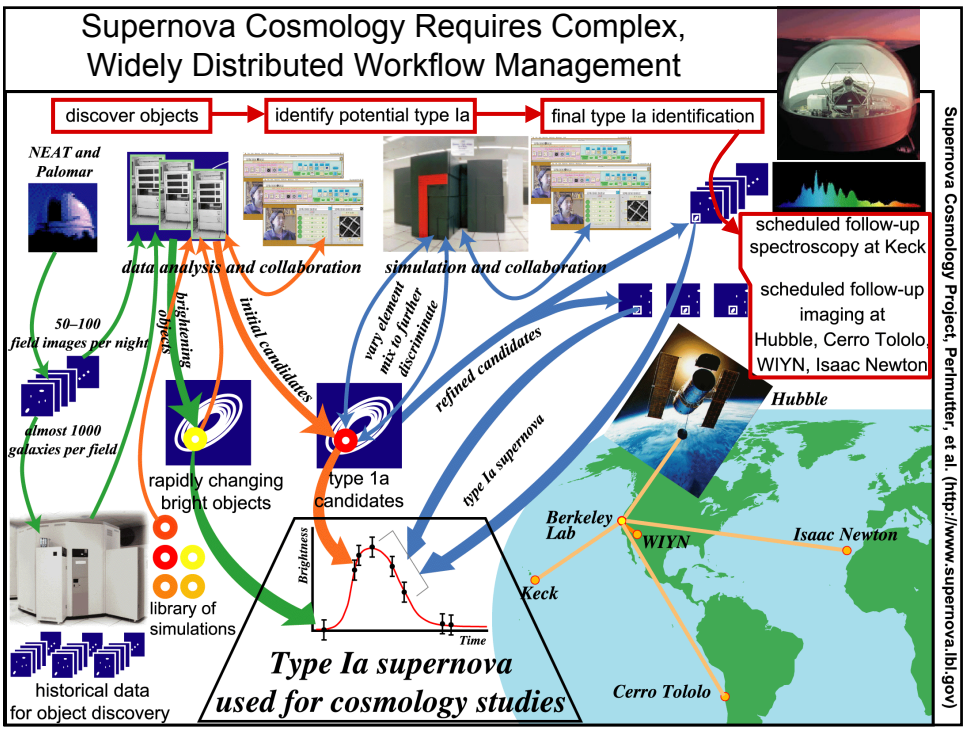

Supernova Cosmology

- This is one of the class of problems that Grids are focused on. It involves:
  - management of *complex workflow*
  - reliable, *wide area, high volume data management*
  - *inclusion of supercomputers in time constrained* scenarios
  - easily accessible pools of computing resources
  - *eventual inclusion of instruments that will be semi-automatically retargeted based on data analysis and simulation*
  - next generation will generate vastly more data (from SNAP - satellite based observation)

What are Grids?

- Middleware for *uniform, secure, and highly capable access to large and small scale computing, data, and instrument systems, all of which are distributed across organizations*
- *Services supporting construction of application frameworks and science portals*
- *Persistent infrastructure for distributed applications* (e.g. security services and resource discovery)
- 200 people working on standards at the IETF-like *Global Grid Forum* (www.gridforum.org)

Grids

- There are several different types of user of Grid services
  - discipline scientists
  - problem solving system / framework / science portal developers
  - computational tool / application writers
  - Grid system managers
  - Grid service builders
- Each of these user communities have somewhat different requirements for Grids, and the Grid services available or under development are trying to address all of these groups

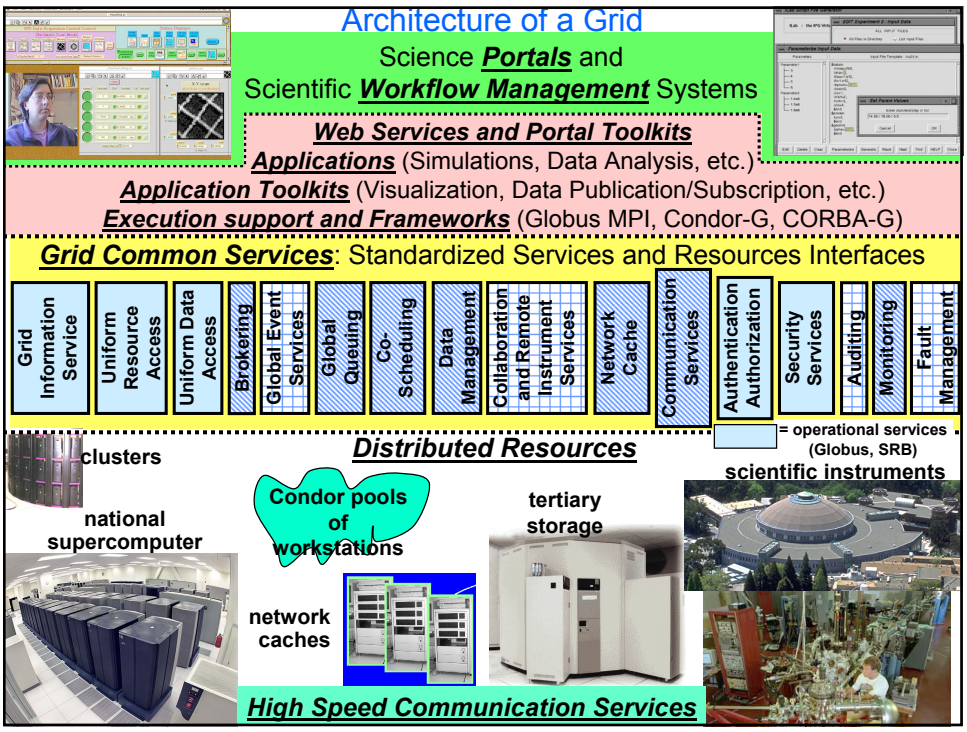

➢State of Grids

- Grids are real, and they are useful now
- *Basic Grid services are being deployed* to support uniform and secure access to computing, data, and instrument systems that are distributed across organizations
- *Grid execution management tools* (e.g. Condor-G) are being deployed
- *Data services*, such as uniform access to tertiary storage systems and global metadata catalogues, are being deployed (e.g. GridFTP and Storage Resource Broker)
- *Web services* supporting application frameworks and science portals are *being prototyped*

State of Grids

- *Persistent infrastructure is being built*
  - *Grid services are being maintained* on the compute and data systems of interest (Grid sysadmin)
  - cryptographic authentication supporting *single sign-on* is provided *through Public Key Infrastructure*
  - *resource discovery services* are being maintained (Grid Information Service – distributed directory service)
  - This is happening, e.g., in the *DOE Science Grid, EU Data Grid, UK eScience Grid, NASA's IPG*, etc.
  - **For DOE science projects, ESNet is running a PKI Certification Authority and assisting with policy issues among DOE Labs and their collaborators**

➢DOE Science Grid

- SciDAC project to explore the issues for providing persistent operational Grid support in the DOE environment: LBNL, NERSC, PNNL, ANL, and ORNL
  - *Initial computing resources*
    - ≈ 10 small, medium, and large clusters
  - High bandwidth connectivity end-to-end (high-speed links from site systems to ESNet gateways)
  - *Storage resources*: four tertiary storage systems (NERSC, PNNL, ANL, and ORNL)
  - *Globus providing the Grid Common Services*
  - Collaboration with ESNet for security and directory services

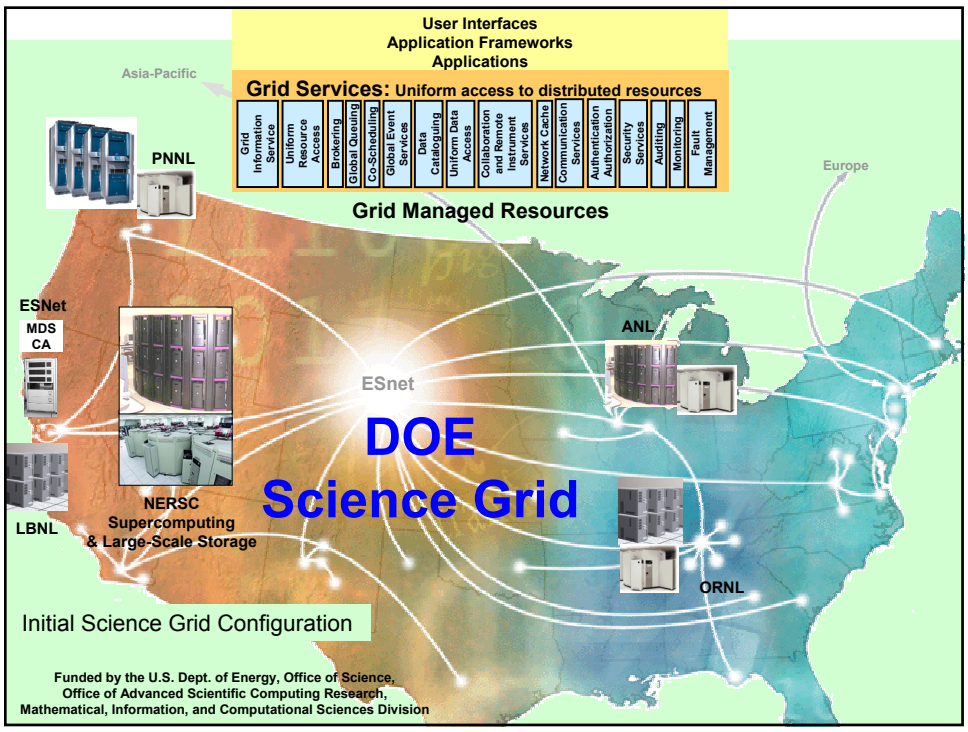

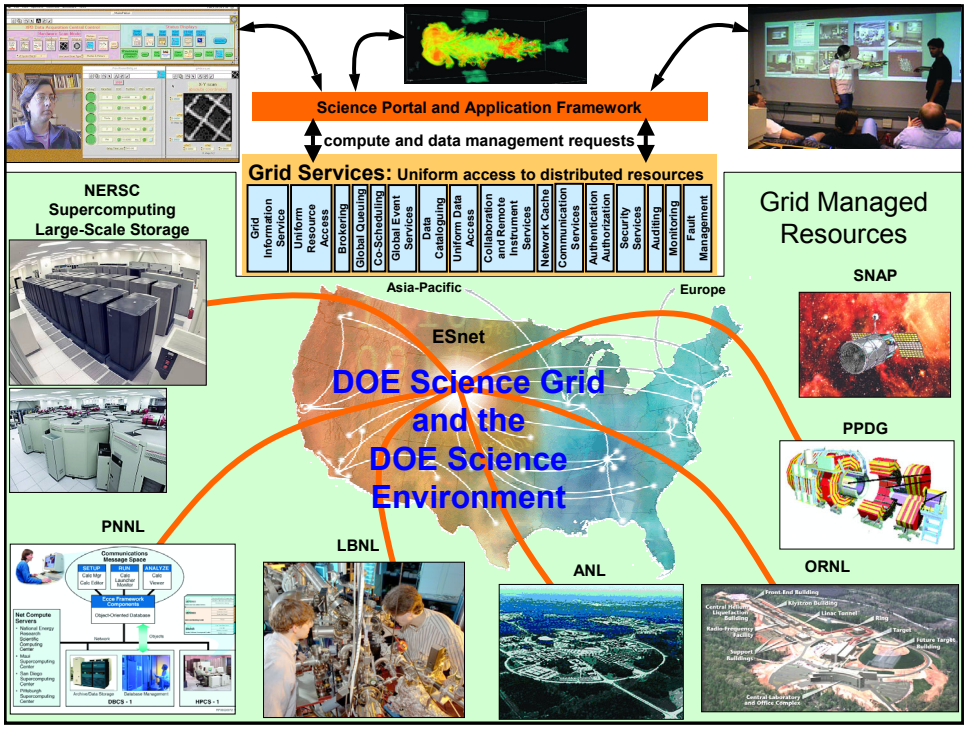

The DOE Science Grid Program: Three Strongly Linked Efforts

- How do we reliably and effectively deploy and operate a DOE Science Grid?
  - Requires coordinated effort by multiple labs
  - **ESNet for directory and certificate services**
  - Manage basic software plus import other R&D work
  - What else? Will see.
- Application partnerships linking computer scientists and application groups
  - How do we exploit Grid infrastructure to facilitate DOE applications?
- Enabling R&D
  - Extending technology base for Grids
  - Packaging Grid software for deployment
  - Developing application toolkits
  - Web services for science portals

Roadmap for the Science Grid

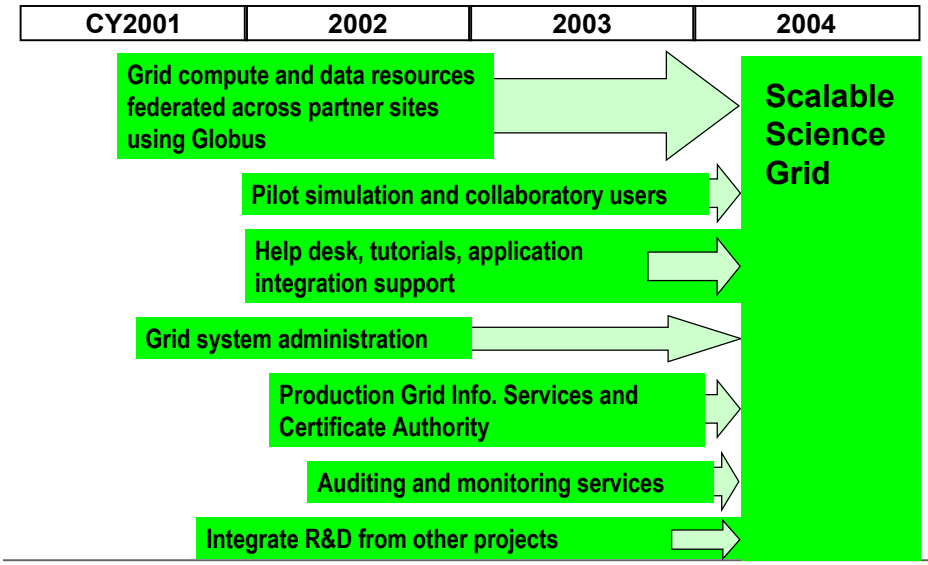

➢ *SciDAC Applications and the DOE Science Grid*

- SciGrid has some computing and storage resources that can be made available to other SciDAC projects
  - By "some" we mean that *usage authorization* models do not change by incorporating a systems into the Grid
  - To compute on individual SciGrid systems you have to negotiate with the owners of those systems
  - However, all of the SciGrid systems have *committed to provide some computing and data resources* to SciGrid users

SciDAC Applications and the DOE Science Grid

- There are several ways to "join" the SciGrid
  - As a *user*
  - As a *new SciGrid site* (incorporating your resources)
- There are *different issues* for users and new SciGrid sites
- Users:
  - *Users will get instruction* on how to access Grid services
  - Users must obtain a SciGrid PKI *identity certificate*
  - There is some *client software* that must run on the user's system

SciDAC Applications and the DOE Science Grid

- New SciGrid sites:
  - New SciGrid sites (*where you wish to incorporate your resources into the SciGrid*) need to join the Engineering Working Group
    - This is where the joint system admin issues are worked out
    - This is where Grid software issues are worked out
    - Keith Jackson chairs the WG

SciDAC Applications and the DOE Science Grid

- New SciGrid sites may use the *Grid Information Services* (resource directory) of an existing site, or may set up their own
- New SciGrid sites may also use their own *PKI Certification Authorities*, however the issuing CAs must have published *policy compatible* with the ESNet CA
  - Entrust CAs will work, in principle – however there is very little practical experience with this, and a little additional software may be necessary

Science Grid: A New Type of Infrastructure

- *Grid services* providing standardized and highly capable distributed access to resources used by a community
- *Persistent services* for distributed applications
- Support for building *science portals*

Vision: *DOE Science Grid will lay the groundwork to support DOE Science applications that require, e.g., distributed collaborations, very large data volumes, unique instruments, and the incorporation of supercomputing resources into these environments*